

# **Strategic Plan for Network and Virtual Services (FY2010)**

Prepared by: David Coder, Ray Pasetes, Mike Rosier, Phil DeMar

Prepared on: 17 July 2009

Last updated: 14 August 2009

## ***Mission***

- Provide all aspects of networking support: architecture, design, acquisition, installation, operation, monitoring, maintenance, and documentation of cabling plant, device infrastructure, and network services necessary to support the Laboratory's network needs.
- Provide centrally managed virtualization services to support the Laboratory's scientific and business needs.
- Provide centrally managed storage and data protection in support of the Laboratory programs. Work closely with our customers to deliver storage solutions that meet or exceed our customer's expectations while maintaining the manageability, flexibility, functionality and security of the services we provide.
- Work with collaborations having special distributed system needs to achieve their performance requirements.
- Engage in network research and development efforts that directly benefit the Laboratory's pursuit of its scientific mission.

## ***Context and Assessment of Current State***

The strategic plan for networking is guided by the strategic plan for Central Services, the strategic plan for Scientific Computing Facilities, the Computing Division Strategic Plan and the strategic directions of the Laboratory; MOU's, SLA's and OLA's established with internal and external organizations; and by Computer Security guidance and directives.

Responsibility for network infrastructure and services at the Laboratory, except for the Accelerator Division's local-area networks, resides with the Computing Division (CD):

- The Network Services (CD/LSCS/NVS/NS) Group is responsible for data forwarding, the physical data cable plant, equipment installation services, network support and services such as DNS, DHCP, NTP.
- The Storage Network Services (CD/LSCS/NVS/SNS) Group is responsible for storage area networking (SAN) and network attached storage (NAS) servers, storage arrays, storage switching and access controls.
- The Virtual Services (CD/LSCS/NVS/VS) Group is responsible for promoting and implementing enterprise-class virtualization services based on a virtual server infrastructure, including host servers, virtual network switching, virtual storage and virtual infrastructure security. The Virtual Services group will build upon the basic virtual infrastructures of the former MIS Department, TD and CD.

- The Wide Area Networking and Network Research (CD/SCF/DMS/WAN) Group is responsible for providing off-site network services and supporting wide-area network infrastructure, including shared management of the ESnet Chicago-area MAN, with neighboring Argonne National Laboratory. The Wide-Area Networking group is also responsible for network infrastructure support of large scale, high performance scientific computing facilities dedicated to high impact distributed movement on a wide-area scale, such as the US-CMS Tier-1 Center. Network research & development associated with scientific computing facilities is also a group responsibility.

### Network Services

The network consists of a mixture of current and obsolete technologies and is reaching capacity limits in several areas:

- Approximately 25% of the network equipment is classified as End-of-Life (EOL) or End-of-Sale (EOS) by the manufacturer;
- Network hubs and 10Base2 Thin-Net coaxial cable is still in production use, though such technology has been largely extinct in networks for a decade;
- Many buildings in the fixed-target areas and Site 37/38/39 areas are connected by aging multimode fiber optic cable. This cable may be capable of supporting 1Gb connections in some areas depending on the condition of the fiber and the distance between switches but is not capable of 1Gb in all areas, nor capable of 10Gb in any area;
- Chassis-based switches are not capable of operating at maximum performance due to the mixture of "Classic", CEF256 and CEF720 modules in the same chassis, plus most 6500-series switches are operating in Hybrid IOS/CatOS mode rather than in Native IOS mode;
- The Wilson Hall network consists of managed and unmanaged switches, and also unmanaged hubs. Fiber optic cables from offices and cubes to the 8<sup>th</sup> floor Fiber Central terminate in obsolete MT-RJ connectors, requiring new network interface modules in Fiber Central switches and new terminations to support gigabit uplinks;
- The DNS system frequently fails and needs to be replaced with an IP Address Management (IPAM) solution to accommodate IPv6 and mandated DNSSEC requirements.
- The DHCP system has reached its capacity and is often overwhelmed by address requests when large numbers of client systems request IP addresses, such as when there are conferences on site.
- Large numbers of high performance computers and storage servers are being added with corresponding saturation of network interfaces. Many switches and interface modules need to be upgraded from 1Gb to 10Gb or from 10Gb to 20Gb or 40Gb with corresponding cable upgrades from multimode to single mode within buildings.

In FY2009, one external assessment of networking was completed. This assessment evaluated the following areas: Network Infrastructure, Critical Network Services, Network Backups, Strategic Plan, Network Management, VoIP, WAN Circuits and Internet Connectivity. This

assessment resulted in a series of action items to improve the performance and reliability of network infrastructure and services.

### Wide Area Networking and Network Research

The current state of the Fermilab wide-area network infrastructure is a highly reliable, high performance, capacious facility operated at the forward edge of network technology. The Chicago metropolitan-area network (MAN) provides the Laboratory with a resilient, flexible, and scalable optical network infrastructure that is capable of satisfying its current and projected off-site bandwidth needs. The Laboratory network perimeter infrastructure is redundant and capacious. The US-CMS Tier-1 Center's network infrastructure has been upgraded to very high performance, high density Cisco Nexus 7000 core switching devices, and is well positioned to accommodate the extremely high local and wide-area network demands of that facility. Network research efforts that align with and directly support the Laboratory's wide-area data movement activities are under way. The major focus area for future enhancements is upgrading the MAN and WAN perimeter infrastructures to be capable of supporting 100Gb/s network technologies.

### Virtual Services

Currently, there are 2 virtual host servers running in a development environment (Wilson Hall 5 West Server Room) and 1 host server running in a production environment (FCC2). All storage is hosted on HP EVA 6100's, which are considered enterprise-class storage arrays. Virtual machines and storage can be moved from one datacenter to another datacenter if configured properly. HP SIM is currently being used to monitor virtual host server hardware. The current virtual infrastructure contains virtual machines, which can reside behind a network firewall or on a DMZ network. This is accomplished by creating separate physical VLANs using interfaces on a network firewall, which in turn connects to multiple switches.

Though there are some virtual desktops running in the environment described above, there should be more consideration to the idea of expanding the use of virtual desktops to areas where sensitive information is accessed on a daily basis or where certain configurations are required.

The potential benefits of running a virtual infrastructure in our current environment are not being fully realized due to lack of resources and higher priority projects. Moving forward with a dedicated group will help the lab use current resources more efficiently, thus minimizing future growth in datacenter space and energy consumption.

### Storage Network Services

Formed in January 2009 as part of a re-organization which was partly influenced by the ITIL/ISO20K project, the dedicated central storage group provides 4 main customer services:

1. SAN – Direct connection to centrally managed disk for block access.

2. NAS – Indirect connection to the centrally managed disk via an IP network for file access.
3. AFS – A global file system accessible from anywhere in the world.
4. Backups – Data protection of lab servers and non-experiment data.

The centrally managed shared resources are used lab-wide, with customer use spanning both scientific and business needs. Currently over 5000 systems use the SAN and NAS services to access ~1PB of storage. The AFS service – a major component underpinning the central web servers and several neutrino experiments – is accessed by collaborators around the world. Demand for the backup service continues to increase ~30-40% annually.

Assessment of current state of each service:

#### Overall

- Monitoring must be improved throughout all services to provide a better understanding of the relationships between the various components. Currently, there is a lack of an end-to-end monitoring system to fully see the performance of all components. Current monitoring is reactionary. Most reports are manual.

#### SAN

- A major upgrade of the SAN switches has placed the fabric (network components only, not storage) in good position. There is ample capacity and the new equipment is well positioned to adapt to upcoming CEE (DCE) technologies.
- There is NOT enough storage capacity on the SAN in terms of performance and space in order to react to changes in production which may adversely affect a shared resource. In 2009, several actions taken to solve performance problems involving FermiGrid, Minos and D0 proved that a higher amount of reserve capacity must be maintained in order to react to (unpredicted) increased workloads from the various experiments. Failure to provide this reserve could result in the shared resource negatively affecting many areas simultaneously.
- Further isolation between experiment areas will need to be achieved to reduce impact. This can be achieved by reducing the sharing of storage arrays between experiments.

#### NAS

- A major upgrade of the NAS servers in 2009 has significantly increased the performance capacity of the NAS service. Additional servers may be needed in the short term depending on the additional load from the upcoming neutrino experiments or increased activity from existing experiments (CDF, D0, FermiGrid, ILC, Minos, Minerva)
- Due to the highly shared nature of the NAS system and the lack of QoS control available from the vendor, access controls must be addressed at the end-user systems.
- Increase education and involvement in customer processes to better understand and develop solutions that meet customer expectations, or at the minimum, set expectations based on funded solution.

## AFS

- Sufficient storage and performance exists to maintain normal growth for the next 2-3 years.
- Automated tools should be developed to reduce the management overhead required when evacuating/re-populating the AFS servers for maintenance tasks.

## Backups

- Though the current backup system is adequate today, it is becoming apparent that improvements are required to:
  - Meet the exploding data growth (adapt to 100-200% annual data change)
  - Protect different types of structured (DB) and unstructured file types
  - Provide encryption
  - Provide archiving
  - The operations effort to maintain the existing backup service is high due to the highly user-interactive nature of resolving backup issues. Many issues are the result of a client system error (misconfiguration, upgrade, etc) that require active involvement from the service maintainers.

## ***Vision***

By 2014, the Laboratory will require local and wide-area network facilities capable of forwarding at hundreds of gigabits per second, storing and accessing over a hundred petabytes of data, and virtualizing hundreds of servers and workstations. Distributed computing collaborations will place significant demand on the networking and storage infrastructures, requiring the ability of network, storage and virtual server systems to handle very high peak network and data processing loads. Network infrastructure will be instrumented to interact with applications & middleware, and adaptable to meet their data movement requirements. End-to-end data circuits will provide wide-area network path bandwidth and latency guarantees for distributed applications.

By 2014, the network will be architected in a modular topology to meet the service level requirements of experiments and users of the Lab's network, storage and virtualization services. Based on MOU's and SLA's, various modules will require 99.999% reliability (no more than approximately 5 minutes of service disruption per year), other modules will require different levels of reliability that can be engineered into each module as appropriate.

By 2014, IPv6 will be fully deployed and operational in local and wide area networks.

By 2014, IP telephony (VoIP) will be fully deployed and operational at the Lab.

By 2014, enterprise-class server and switch virtualization services will be fully deployed and operational at the Lab. The use of virtual infrastructures will be just as common as stand-alone servers for most general-purpose computing applications.

By 2014, enterprise storage will consist of a back-end storage infrastructure with a user-facing service of data availability and protection. The exact back-end technology will be based on MOU, SLA or OLA criteria of storage space, access speed, security (encryption), access frequency (for hierarchical storage capabilities), and long-term archival requirements. SLA's will include Disaster Recovery or Business Continuity criteria (Recovery Time Objective and Recovery State Objective).

## ***Stakeholders***

The list of stakeholders includes everyone at the Lab who uses networking, shared storage or virtual servers, including:

- Experiments such as CDF, DZero, MINOS, MiniBooNE, MINERvA, and Grid collaborators;
- Off-site experiments such as NOVA, DUSEL, DSM, MINOS, CDMS, and Auger;
- Providers of wide-area and Internet services to the Lab such as ESnet;
- Experiments with specific service requirements for data movement, storage, performance or reliability such as CMS and the LHC Remote Operations Center (ROC).

## ***Strategic Goals***

The strategic goals for networking are based on a set of high level, architectural principles. These principles define a core philosophy that helps ensure decisions on the design, implementation, and upgrade of the Laboratory network are made consistent with a common strategic direction. They should provide the basis for all levels of networking decisions within the organization, from the design of major projects to the implementation of small project tasks. The strategic goals are not drafted to be specifically applicable to any one component of this plan.

### **Shared Goals of Network Services & Wide-Area Networking Groups**

- Network designs or configurations will be kept as simple as requirements allow in order to minimize implementation and support effort, improve reliability, and simplify troubleshooting by following established industry Best Practices such as the Cisco Enterprise Campus Architecture, the Cisco Enterprise Composite Network Model and the Cisco Enterprise Data Center Architecture;
  - Modularity can help make each design element simple and more easily understood; can establish failure domains, security policy domains, availability/performance domains, etc.
  - Modularity can simplify network scalability by scaling within and between modules.
  - Modular designs allow the addition or removal of modules as new services and functions are needed without changing the underlying network design.
- Network infrastructure capacity needs to be significantly increased in order to meet present MOU's and SLA's;

- High capacity, high density switch fabric will be used to the greatest extent practicable, in order to minimize management and maximize performance;
- Network infrastructure will be restored to the forward edge of established network technology, while following industry Best Practices of architecture and design;
- Network reliability needs to be improved through the use of redundancy where appropriate, or by alternative routing and switching paths where appropriate.
- Users and associated computing resources affiliated with a specific organization, experiment, project or location should be aggregated into network modules in order to provide structure to the facility network. A modular topology will allow the implementation of technology and reliability appropriate to meet the SLA's of each module;
- By default, data forwarding path determination will be via the use of standard routing protocols, specifically OSPF and BGP, with use of policy or other types of routing limited to alternate path forwarding needs for specific applications or facilities.
- Network support staff requires training to keep up with current network technologies and methodologies.
- Commercial service offerings and supported open-source tools & software will be preferred over in-house solutions.
- Network monitoring infrastructure will be deployed consistent with a vision of providing network information to users, applications, and middleware.

#### Network Services Group Specific Goals:

- Wireless network media will be deployed as an integral component of general network access;
- Remote access from off-site locations via Virtual Private Networking (VPN) will be deployed as an integral component of general network access;
- Problematic and obsolete technology will be replaced by current, more flexible and more reliable products, for example:
  - The QIP DNS/DHCP system will be replaced with an appliance-based system that is capable of scalability and DOE-mandated security (DNSSEC);
  - The Cisco ACE and Alteon load balancers will be replaced by an F5 application delivery appliance;
  - The GPS receivers for network time (NTP) will be replaced with current, supported GPS and CDMA receivers;
  - End-of-Life/End-of-Sale (EOL/EOS) equipment as designated by the manufacturer, and obsolete equipment such as hubs.

#### Wide Area Networking & Network Research Group Specific Goals

- Redundancy designed into the Laboratory's off-site network access will be based on complete physical and geographical diversity, to the maximum extent practical.

- Aggregate off-site network capacity will continuously be maintained at a level to meet peak traffic load requirements for Laboratory facilities, experiments, and applications.
- Performance analysis services for Laboratory wide-area, distributed applications will be automated to the extent practicable.
- Research efforts in network monitoring (perfSONAR) and data circuit services will be integrated into and directly support Laboratory facilities, experiments, & applications.

#### Virtual Services Goals

- Use virtualization as a driver toward more centralized IT services. Virtualization requires a centralized environment to reach the economies of scale required for maximum cost savings.
- Provide virtualization services that have broad customer appeal by aggressively communicating the many benefits that can be achieved through the use of virtualization technologies.
- Provide a virtual infrastructure which provides rapid provisioning of virtual machines and applications.
- Provide cost-savings analysis data to support increased use of virtualization technologies at Fermilab.

#### Storage Network Services Goals

- Flexibility -- the service should stay nimble enough to adapt the services to changing workloads while minimizing service outages.
- Maintain enough capacity in terms of storage space and performance in order to support the above objective.
- Keep staff trained in the relevant areas for the specific technologies placed into production.
- Investigate new technologies which can add to or enhance current services.
- Implement monitoring to aid in diagnostics, provide metrics and trending information as well as show the value of the various services to our customers.
- Architectural designs should take into account virtual server and virtual networking technologies. Where appropriate, couple solutions around these technologies to provide flexibility, simplicity, manageability and security.
- Automate repetitive tasks as much as possible; provide documentation and FAQs to reduce operation effort as much as possible.
- Engage users to better understand requirements, use cases and expectations. Educate customers in the capabilities of the services, document known limitations and design considerations to aid customers in the use of the services.

### ***Strategic Objectives:***

Strategic objectives are tangible targets for efforts or activity areas that are intended to be the means of achieving strategic goals. They may be specific enough to be applicable to only one

area of activity, or may span multiple areas. There are normally timeframes associated with strategic objectives.

### Network Services Objectives

- Upgrade core network capacity and reliability to meet or exceed MOU, SLA and OLA requirements:
  - 2010 – Upgrade core switch fabric to 2Tb capacity with 80Gb backplane capacity;
  - 2010 – Start to deploy MPLS in the network core;
  - 2011/2012 – Deploy next generation (40/100Gb) backbone links as needed and as available;
- Upgrade building/experiment distribution switch capacity and reliability to exceed MOU, SLA and OLA requirements:
  - 2010 – Upgrade distribution switch fabric to 2Tb capacity as necessary;
- Upgrade or replace obsolete network devices as necessary to assure network performance and reliability. This includes devices for which software or firmware upgrades are no longer available, which use obsolete data link technologies, or are identified as EOL/EOS by the manufacturer, or which don't support SSH:
  - 2010 – Upgrade all Hybrid CatOS/IOS switches to Native IOS software;
  - 2010 – Replace all CatOS-only switches with IOS equivalents;
  - 2010 – Replace all shared-media devices (e.g. hubs) with dedicated media devices (switches).
  - 2010/2011 – Replace remaining Category 3 cable with Category 5e or 6 as appropriate.
  - 2011 – Improve chassis switch performance by replacing "classic" bus modules with CEF720 or better modules;
  - 2011 – Upgrade access switches to support 1000BT;
  - 2011 – Upgrade/replace approximately 50% of the remaining EOL/EOS devices;
  - 2012 – Upgrade/replace remaining EOL/EOS devices;
- Upgrade DNS/DHCP application
  - 2010 – Implement appliance-based IP address management system, replacing the error-prone QIP system;
  - 2010 – Implement DNSSEC as mandated by the DOE;
- Upgrade Wireless networks
  - 2010 – Upgrade wireless access points in FCC and Wilson Hall
  - 2010 – Install wireless mesh network in Village and selected Fixed Target areas
  - 2011 – Expand wireless coverage to all areas of the Lab as appropriate
  - 2012 – Upgrade WiSM modules to next-generation Wireless LAN Controllers
- Upgrade physical cable plant
  - 2010 – Install new fiber-optic cable between Wilson Hall and LCC (for GCC) to support new communication requirements and to provide path diversity between FCC and GCC.

- 2010 – Upgrade fiber and copper cable infrastructure in FCC
- 2010/2011 – Replace multimode fiber between buildings in the Site 38 area with single mode fiber to support higher throughput.
- Implement Network Access Control system
  - 2010 – Deploy Cisco NAC solution to authenticate users, enforce minimum security criteria to connect to the network, and to help identify the types of computers and other devices connected to the network.
- Upgrade network security systems
  - 2010 – Implement network firewall systems in Wilson Hall and FCC distribution switches.
  - 2010 – Implement network firewall systems for Visitor wireless connections.
- Implement IP Telephony (VoIP)
  - 2010 – Upgrade the VoIP pilot infrastructure to support Exchange integration for Unified Communication, and to support scalability to be able to expand VoIP usage across all areas of the Lab.

#### Wide Area Networking and Network Research Objectives

- Off-site network device infrastructure and operations
  - Facilitate, support, and upgrade as necessary, a fully redundant perimeter network infrastructure that provides the Laboratory with the high bandwidth data channels necessary for its offsite data movement requirements. Timelines:
    - 2010/2012 - Increase off-site 10GE data channels to meet data movement requirements
    - 2010/2011 - Deploy next generation end-to-end circuit router & border routers
    - 2011/2014 - Implement 100Gb/s off-site links on perimeter infrastructure
  - Deploy capacious end-to-end data paths to remote sites involved in high impact data movement with the Laboratory, including support infrastructure to facilitate use, monitoring, and troubleshooting of those paths. Timelines:
    - 2010/2014 - Facilitate optimal use of network paths between CMS T0/T1/T2 sites, including increasing bandwidth capacity to meet evolving requirements, and deployment of dynamic circuit capabilities
- Network support & analysis services for distributed, high-impact scientific computing:
  - Deploy higher density, higher capacity network infrastructure in support of high impact scientific computing facilities. At the current time, this is limited in scope to the CMS Tier-1 facility. Timelines:
    - 2010 - Complete upgrade of CMS Tier-1 facility LAN to be redundant core 10GE-based data center switches, with very high capacity switch fabric
    - 2010/2011 - Migrate Tier-1 servers to direct connections on CMS data center switches; support initial implementation of 10GE-connected CMS host systems

- 2010/2012 - Upgrade interconnections to 100Gb/s or more between core switches, and to 80Gb/s for access switches
  - 2012/2014 - Expand data center switches to high density chassis, higher capacity switch fabric; replace access switches with data center-class switches
- ESnet Chicago Metropolitan Area Network (MAN) design, engineering, and operations
  - Operate the Chicago MAN, which supports all of the Laboratory's off-site access, at a very high level of reliability, increase MAN channel deployment to meet bandwidth needs, introduce next generation optical network technology as it emerges, and develop a MAN test-bed capability. Timelines:
    - 2010 - Establish 2nd Ciena DWDM hub on site, geographically separate from FCC
    - 2010 - Implement a 10GE channel to ANL to serve as an initial MAN testbed
    - 2010 – Upgrade the west leg (FNAL-ANL-Level3 PoP) of the MAN to 20GBS SDN capacity, for closer parity with the east leg
    - 2011/2014 - Upgrade the MAN to 100Gb/s capability & deploy 100GE channels
    - 2013/2014 - Replace or renew the ComEd fiber component of the MAN, as appropriate
- Network research & development efforts
  - Develop an R&D support infrastructure to participate in international & national, advanced, optical wide area network research initiatives. Timelines:
    - 2010 - Deploy a full service IPv6 test-bed network facility
    - 2010 – Develop content management system for performance analysis
    - 2010/2012 - Develop & deploy End Site Control Plane systems
    - 2010/2012 - Develop Network Weather & Performance Service E-Center
    - 2011/2012 – Automate performance analysis methodology
    - 2010/2014 - Support Laboratory participation in additional advanced network and distributed systems R&D projects and collaborations

### Virtual Services Objectives

- Develop a long-range plan to deal with future architecture, budget, scope, staffing, and technology selection concerns:
  - 2010 – Meet with CD management and stakeholders to gather ideas and expectations, which will help set the priorities and direction of the virtual services group.
  - 2010/2011 – Develop documentation to describe the current and future virtual architecture, including networking, security, failover, storage, dependencies, and other key elements to aid in long-range planning.
- Upgrade current MIS virtual infrastructure to support additional production virtual machines:
  - 2010 – Build out MIS virtual infrastructure to by adding 6 additional CPU's,

- memory, and licenses across 2 VMware ESX servers.
  - 2010 – Expand SAN storage capacity with new drives to support additional virtual machines.
  - 2010 – Install a second VMware vCenter Server to provide redundancy of management and monitoring tools.
  - 2010/2011 – Upgrade MIS virtual infrastructure to vSphere 4 Enterprise Plus version.
- Implement consolidated backups in virtual environment:
  - 2010/2011 – Investigate, design, and implement a backup strategy for virtual machines utilizing VMware’s consolidated backup feature.
- Implement additional capabilities in analysis, monitoring, and reporting of metrics in a virtual environment:
  - 2010 – Investigate and Implement 3<sup>rd</sup>-party or built-in tools used to monitor production environments.
- Demonstrate viability of new application delivery mechanisms for production use:
  - 2010/2011 – Investigate/Install/Configure VMware View ThinApp , Microsoft App-V, or other application virtualization products.
  - 2012 – Investigate and install load-balanced Terminal Servers, TS Remote App, or other applications used for running applications using terminal services.
- Demonstrate viability of thin client/virtual lab technologies for production deployment:
  - 2011 – Install and configure thin clients for use as kiosks, training machines, and contractor machines.
  - 2012 – Incorporate technologies that allow fast provisioning and reimaging of virtual desktops.
- Implement a tiered services architecture utilizing various virtualization technologies working closely with storage and networking groups:
  - 2011 – Compare virtualization offerings by multiple providers to identify the most cost-effective products for each type of virtualization technology in use.
  - 2012 – Identify each tier (i.e. desktops running on iSCSI storage, servers running on Fiber Channel storage, highly available servers running on virtualized storage) and select the software, hardware, protocols, and other components which will be used to implement each one.
  - 2012 – Develop processes and tools to aid customers in selecting a tier that meets their performance needs and budget constraints.

#### Storage Network Services Objectives

- Service Consolidation
  - SAN
    - 2010 Merge MIS/CD/(TD) SAN Networks – includes incorporating existing SAN fabrics in WH/CD via FC and possible iSCSI SAN at TD
    - 2011 Replace aging equipment in MIS/TD with standardized SAN equipment (HDS)

- Backups
  - 2010 Merge MIS/CD/ (TD) Backups – investigate existing requirements and match to existing infrastructure. Where required, upgrade existing infrastructure to meet/exceed requirements from MIS (and TD).
  - 2010 Implement hardware based encryption services to protect PII data
  - 2010 Implement offsite storage practices w/ 10yr retention. Includes the planning of keeping proper equipment to actually recover data that is 10yrs old and developing a plan to test periodic restoration and verification of data.
- Service Upgrades
  - AFS
    - 2010 Upgrade AFS servers to v1.6 of software using Solaris ZFS filesystem.
    - 2011 Replace AFS servers with Linux servers.
    - 2012 Upgrade AFS service with new features from YFS SBIR project.
    - 2013 Integration of mobile devices (Windows Mobile, Palm, etc) with AFS/YFS file system for data access.
  - Backups
    - 2010 Implementation of disk based library for TiBS, further increasing per server backup capacity.
    - 2010 Expansion of Commvault backup system to fully support Exchange and SharePoint services for both backups and archiving.
    - 2010 Increase LTO-4 capacity as appropriate for service demand.
    - 2010 Continue investigation into Single-instance-store (data de-duplication) and Virtual tape libraries for feasibility and price point that is appropriate for the applications at the lab.
  - NAS
    - 2010 Investigate/implement the use of BlueArc Mercury platforms. These new platforms are a low-cost alternative to existing Titan infrastructure that also underpins new NFSv4.1 server implementation. This new implementation allows for multiple servers to serve data in parallel providing better scaling for the shared NAS service.
    - 2011 Implementation of NFSv4.1 pNFS with integration into Kerberos infrastructure.
    - 2011 Investigation into Tier 0 storage (SSD) to be used in conjunction with existing storage devices to provide better performance and scalability.
    - 2012 Investigation of incorporating an HSM system into the NAS service
    - 2013 Implementation of HSM system into NAS service
  - SAN Fabric
    - 2010 Complete upgrade of SAN Fabric to 4Gb.
    - 2011 Implement CEE (DCE) and upgrade to 8Gb fibre w/ 10Gb IP uplink.
    - 2012 Implement full CEE with existing 40Gb/100Gb IP infrastructure
    - 2013 Implement DR/BC

- SAN Storage
  - 2010 Replace aging HDS storage with 2<sup>nd</sup> generation storage
  - 2010 Replace aging 3PAR storage (reaching EOS) with 2<sup>nd</sup> generation storage
  - 2010 Expand existing storage capacity and performance to fully support Exchange, SharePoint, Virtualization Service, migration of DB servers to SAN and expansion of AFS, Backup and NAS services.
  - 2010 Research/implement DDN S2A6620 storage devices. If proven to work, these devices will provide a lower cost storage alternative for experimenters that will work with the NAS service at a price point that is sustainable and ensures experiment isolation.
  - 2010 Investigate the use of hardware virtual controllers (USP-V) or s/w based virtual controllers (such as Falconstor) to add flexibility and reduce downtimes required for storage reconfigurations/migrations.
  - 2011 Implementation of Virtual storage controllers
  - 2011 Research/Investigate costs and requirements for disaster recovery and business continuity (DR/BC).
  - 2012-2013 Implement DR/BC for storage.
- Service Operations
  - 2010 Establishment of user meetings to promote services, better understand user needs and issues and raise awareness of the limits of the various services and use cases.
  - 2010 Implement SAN monitoring that is heterogeneous, aids in diagnostics, provides trending as well as modeling capabilities.
  - 2010 Provide adequate training to personnel in new technologies/equipment used in production.
  - 2010 Increase documentation using established document repository (docdb or SharePoint) for user training, document processes, etc.
  - 2011 Establish processes and procedures that will comply with the upcoming ITIL/ISO20K requirements for Change Management and Continuous Improvement Plans

## ***Strategies***

### Network Services

The strategy for network architecture and networking support is the Fermilab Modular Architecture, a site-specific implementation of the Cisco Enterprise Campus Architecture and the Cisco Enterprise Composite Network Model in order to improve scalability, flexibility and reliability.

- Modularity can help make each design element simple and more easily understood and can establish failure domains, security policy domains, availability/performance domains, etc.
- Modularity can simplify network scalability by scaling within and between modules.
- Modular designs allow the addition or removal of modules as new services and functions are needed without changing the underlying network design.

Fermilab computer rooms are being re-engineered as a site-specific implementation of the Cisco Enterprise Data Center Architecture.

While our model is based on Cisco's best practices architecture, there is no requirement to use Cisco products so we are not "locked in" to any vendor.

### Wide Area Networking and Network Research

High-level wide-area networking strategies employed to achieve our strategic objectives in a manner consistent with our strategic goals include:

- Maintaining and upgrading wide area network-supported infrastructure to keep it as capacious as practical. Capacious network infrastructure helps avoid application-level performance problems, and provides the necessary agility to accommodate changing needs.
- Aligning wide-area networking infrastructure and support with DOE science research efforts, as well as technology trends in the general research & education community. Technology decisions and deployments need to be made in context of the global collaborations the Laboratory supports.
- Integrating network research efforts with the Laboratory's advanced scientific computing activities. Applied network research directly applicable to Laboratory distributing computing efforts benefits both the results of the research, and the Laboratory's overall scientific mission.

### Virtual Services

- Use automation to achieve goals in providing provisioning (i.e. Windows PowerShell or visual basic scripts that interact with VMware API's).
- Implement application virtualization technologies (i.e. VMware ThinApp or Microsoft App-V).
- Expand current virtual infrastructure by adding additional storage and processing.
- Investigate 3<sup>rd</sup>-party applications which enhance the advanced features of our current platform.
- Actively encourage the use of virtualization services, which requires having an environment that is stable, flexible, secure, and easy to use.
- Engage in long-range planning exercises to help set the future direction of virtualization technologies at Fermilab.

## Storage Network Services

- Implement high-availability/redundancy in all areas. Service outages should eventually become an anomaly and not a regularly scheduled maintenance window.
- Implement standards of operation – document processes and procedures for daily operations and changes. Use this documentation to review processes to aid continuous improvement plans (CIP).
- Engage customers early and often to understand needs, issues and better anticipate future requirements.
- Standardize monitoring. Along with alerts, the monitoring should include end-to-end quality of service (QoS) information, trending information and, if possible, modeling information. The latter two can be used for Capacity Management. The monitoring tools should be able to monitor heterogeneous systems. From the tools a set of requirements can be placed on the types of equipment we will choose to purchase.
- Continuous pursuit of new technologies and automation that can reduce operational costs, increase efficiencies and/or change the “way we do things” in a positive manner.
- Maintain a skilled and cross-trained staff to ensure that data protection is not just about the different technologies put into place, but also the personnel required to operate and maintain the systems. The loss of a single person should not place the service in jeopardy.
- Delivered solutions should be modular, flexible, simple, and secure whenever possible. Manageability and maintainability of the solution should be stressed at all levels. Designs should always be cognizant of virtual server and network technologies and deployments. When possible, designs should encompass all three aspects providing a single packaged solution.

## **Resource Needs**

### Network Services

Historically, the level of effort, in terms of both personnel and M&S costs, has remained relatively constant for network support over the years, however significant investment will be required in the next several years:

- To replace the aging fiber infrastructure in several areas of the Lab;
- To replace obsolete equipment presently in service;
- To increase capacity and reliability to meet the increasing needs of experiments and general computing.

### Wide Area Networking and Network Research

Wide-area networking research & development efforts will be increasing as we initiate two new projects funded by the Department of Energy. Part of the grant funding will be used to bring in

additional software development resources for those projects. The rest of the effort for those projects will come out of the existing staff. In most cases, the new R&D effort aligns with current staff job efforts & responsibilities. But some redirection of effort away from current ongoing wide-area network support will be unavoidable. We will require some increases in efficiency and productivity to sustain the Laboratory's wide-area network service at current performance levels. New management tools will need to be procured if that expanding scope of work is to be met within current staffing levels.

Additional effort and skill sets will likely be needed to provide guidance in optimizing use of emerging distributed applications across our increasingly capacious network infrastructure. We anticipate this to become an increasingly important aspect of Laboratory network support for our collaborations that are involved with wide-area distributed computing systems.

Over the longer term, the Laboratory's networking efforts need to be consistent with the DOE Office of Science strategic directions to develop Petascale science facilities supported by terabit/sec network infrastructure. Additional resources, both personnel and capital expenditure, may be needed to keep the Laboratory at the forward edge of the Office of Science's strategic direction.

### Virtual Services

In order to build up our virtual infrastructure in the near future, resources from other groups will likely be required to assist in migrating physical machines to virtual machines. Cooperation with networking, storage, operating system, and application groups will be essential in achieving a successful virtual environment.

Depending on the goals set by management, additional staff or contractors may be necessary to achieve these goals.

### Storage

With respect to the scientific needs, often times the exact requirements are not well known due to the very nature of their business – one of a kind, never been developed. As such workloads can vary significantly from initial prototyping/modeling to actual production and may not have been accurately predicted when initially providing a solution. Sufficient additional capacity in terms of both storage space and storage performance must be maintained in order to allow the services to adapt to changing workloads and patterns of the various customers.

As demand for the various services increases, limited staff resources will negatively impact the following areas in order:

1. Effort to investigate new technologies
2. Effort to investigate improving current services

3. Effort to maintain standard operations

## ***Progress Indicators***

### Network Services

The level of progress in attaining strategic objectives for Network Services will be determined through a combination of four factors:

1. Meeting externally-imposed milestones, such as DOE-mandated policy compliance deadlines;
2. Measurement and publication of performance and reliability data to demonstrate that MOU, OLA and SLA requirements have been met;
3. Service desk ticket statistics showing failure recurrence, time to close, and other statistics to show that the network is becoming more reliable and that the Network Services support staff is responding and resolving issues as required;
4. Feedback from stakeholders.

### Wide Area Networking and Network Research

The level of progress in attaining strategic objectives for this plan will be determined through a combination of three factors:

1. Comparison between the timeline expectations for strategic objectives listed in this plan, and what is actually achieved within those time frames. This comparison is not intended to be absolute. It is expected that there will be some time shifting in implementation of identified objectives, given the dependencies on technological evolution, personnel resources, and changing requirements. Rather, the progress is better gauged by how closely implementation compares to the general trend outlined for the objective.
2. Measurement and observation of how on the capacity and capabilities of the network infrastructure and services compare to the utilization and performance at any particular time. Insufficient capacity or capabilities to meet current requirements is a potential indicator that progress needs to be greater.
3. Feedback from stakeholders. In the end, network infrastructure and support services exist to satisfy the needs of the stakeholders, and they should be the ones to determine how well their needs are being met.

### Virtual Services

The level of progress in attaining strategic objectives for Virtual Services will be determined through a combination of four factors:

1. The number of active virtual servers in production use;
2. Measurement and publication of performance and reliability data to demonstrate that

- MOU, OLA and SLA requirements have been met;
3. Service desk ticket statistics showing failure recurrence, time to close, and other statistics to demonstrate the reliability of the virtual infrastructure and that the Virtual Services support staff is responding and resolving issues as required;
  4. Feedback from stakeholders.

### Storage Network Services

The level of progress in attaining strategic objectives for Storage Network Services will be determined through a combination of four factors:

1. The amount of networked storage in use;
2. Measurement and publication of performance and reliability data to demonstrate that MOU, OLA and SLA requirements have been met;
3. Service desk ticket statistics showing failure recurrence, time to close, and other statistics to demonstrate the reliability of the storage infrastructure and that the Storage Network Services support staff is responding and resolving issues as required;
4. Feedback from stakeholders.

### ***Additional Information***

<text... if any>